Journal of Nuclear Materials 408 (2011) 30–39

Contents lists available at ScienceDirect

Journal of Nuclear Materials

journal homepage: www.elsevier.com/locate/jnucmat

Prediction of radiation induced hardening of reactor pressure vessel steels using artificial neural networks

N. Castin [a,b,*], L. Malerba [a], R. Chaouadi [a]

[a] Studie Centrum voor Kerneenergie – Centre d'Etudes de l'énergie Nucléaire (SCK•CEN), NMS unit, Boeretang 200, B2400, Mol, Belgium
[b] Université Libre de Bruxelles (ULB), Physique des Solides Irradiés et des Nanostructures (PSIN), Boulevard du Triomphe CP234, 1050 Brussels, Belgium

ARTICLE INFO

Article history:
Received 6 September 2010
Accepted 22 October 2010

ABSTRACT

In this paper, we use an artificial neural network approach to obtain predictions of neutron irradiation induced hardening, more precisely of the change in the yield stress, for reactor pressure vessel steels of pressurized water nuclear reactors. Different training algorithms are proposed and compared, with the goal of identifying the best procedure to follow depending on the needs of the user. The numerical importance of some input variables is also studied. Very accurate numerical regressions are obtained, by taking only four input variables into account: neutron fluence, irradiation temperature, and chemical composition (Cu and Ni content). Accurate extrapolations in term of neutron fluence are obtained.

1. Introduction

It is well known that reactor pressure vessel (RPV) steels used in light water nuclear reactors harden and embrittle under neutron irradiation. Embrittlement is customarily measured in terms of increase of the ductile-to-brittle transition temperature (DBTT), measured by means of Charpy tests. Nuclear regulations impose safety margins on this increase, according to rules that may somewhat change depending on the country, as safeguard against RPV failure in both service and accidental conditions [1,2]. Radiation embrittlement of materials depends *a priori* on many variables: not only neutron fluence, flux, and energy spectrum, but also irradiation temperature, chemical composition, and pre-irradiation material history [3]. All these variables must be simultaneously considered to reliably predict pressure vessel embrittlement. However, in order to be able to assess the effect of the different variables on the mechanical response of the steels, each of them should be varied independently of the others in a sufficiently wide range. Such an approach is clearly unrealistic.

Although inadequate to cover all possible conditions, a large amount of data from surveillance capsules and from material test reactors does exist. One of the most important goals for utilities and other nuclear stakeholders is the development, based on "clever" interpolations and extrapolations of the available data, of reliable trend curves, providing estimates of steel embrittlement as a function of the most important among the above-mentioned variables [2].

Artificial intelligence is the combination of algorithms, data and software used to develop computer systems that can be defined *intelligent*. One defining feature of intelligence is the capability of learning from past experience and solving problems when important information is missing, so as to be able to handle complex situations and to correctly react to new circumstances. There are many different computational models which are considered branches of the artificial intelligence field, each one suited to a different kind of problem. Artificial neural networks (ANN), for instance, provide a general framework for representing non-linear functional mappings between a set of input variables and a set of output functions [4]. The list of successful applications of ANN to real-life problems is endless, in sundry domains of interests, e.g. character and image recognition, image compression, stock market prediction, tumor detection in medical image analysis, vehicle piloting, etc. The interested reader can find general information for example in [5,6].

In the field of nuclear materials, ANN have been applied by Kemp et al. [7] to the analysis and prediction of the yield strength increase ($\Delta\sigma_Y$) induced by irradiation in low activation ferritic/martensitic steels, which are candidate structural materials for future nuclear fusion reactors. Their conclusion was very encouraging concerning the ANN capability of analyzing irradiation damage, at least within the range of irradiation parameters and steel composition that are covered in the database used for training. Later on, Windsor et al. [8] have shown that the network can also be used for extrapolating to fluences higher than those included in the training database. In this work, we use a different ANN approach to construct a mathematical regression of the radiation-induced $\Delta\sigma_Y$, as a function of irradiation parameters and steel composition. Differently from [7,8], we compare two different

* Corresponding author at: Studie Centrum voor Kerneenergie – Centre d'Etudes de l'énergie Nucléaire (SCK•CEN), NMS unit, Boeretang 200, B2400 Mol, Belgium.
E-mail address: ncastin@sckcen.be (N. Castin).

0022-3115/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jnucmat.2010.10.039

ANN training approaches (classical and Bayesian), and try to identify the most suitable, depending on the purpose of the trained ANN. Moreover, we define and compare two different algorithms to split the available data in training and validation sets, because this aspect of the ANN training problem is very important for applications where the available amount of training data is limited.

The objectives of this work are manifold:

1. To exploit, as effectively as possible, the information contained in the available databases from surveillance and material test reactors irradiations, for both steels and alloys.
2. To identify in a systematic way the variables that appear to be of higher or lesser importance, based on the available data, within the ranges covered, i.e. based on interpolations. For example, the possible existence of a flux effect is addressed.
3. To attempt an extrapolation outside the ranges covered by the databases and evaluate the reliability of these extrapolations, by assessing the capability of the ANN to predict a certain category of data when trained on a different category. For example, prediction of the evolution of hardening for higher fluences will be attempted.
4. To provide a guide to design future irradiation experiments on steels and alloys, in order to better understand specific effects and dependencies.

Fully reliable predictions will only be possible once the important physical mechanisms acting during irradiation have been identified, understood and quantified at all relevant scales, from the atomic scale to the component scale. However, an empirical approach based on advanced regression techniques, such as ANN, can be beneficial for industrial applications within a shorter delay and can even be useful to guide the longer term development of physical models.

In Section 2, we describe in detail the different ANN training approaches that we propose, and provide some theoretical background for the reader unfamiliar with these techniques. Then, we briefly describe, in Section 3, the RADAMO database that is used throughout this paper. In Section 4, training experiments are reported, aimed at identifying which, according to our ANN, are the most influential factors for hardening. Finally, in Section 5 we compare the different algorithms proposed in Section 2 in order to establish which combination is the most suitable for extrapolation under given conditions.

2. Methodology

Artificial neural networks (ANN) are powerful computational models, capable of providing efficient numerical regressions even when many input variables are involved. In this work, we use the classical feed-forwards multi-layer perception [4] with one hidden layer, linear combination functions and hyperbolic tangent activation functions, as depicted in Fig. 1. It is a *universal approximator* in the sense that it can approximate any continuous multivariate function to any desired degree of accuracy, provided that enough hidden nodes are available [9,10]. The *universal approximation theorem*, however, does not provide a theoretical framework for training ANN, but only demonstrates the existence of at least one ideal architecture for any regression problem, without guaranteeing that it can be found by training, and without giving an estimate of the number of training examples that must be provided.

The output of the ANN shown in Fig. 1 can be written as:

$$\Delta\sigma_Y = \tanh\left(w_{O0} + \sum_{j=1}^{H} w_{Oj} \tanh\left(w_{j0} + \sum_{k=1}^{4} w_{jk} \cdot i_k\right)\right) \quad (1)$$

where H is the number of hidden nodes, i are the input variables. The fixed coefficients w_{O0}, w_{Oj}, w_{jO} *and* w_{jk} are the *synaptic weights* (also often called *synapses*). The advantage of the ANN method is that this generic expression does not require the user to explicitly state how input variables and outputs are related to each others, unlike the usual trend curves [11,12]. All input variables are connected to all hidden nodes, and these play the role of simple processing units that, connected in network, can reproduce complex mappings that are not necessarily visible to the user. The drawback, however, is that no reasonable physical interpretation of the individual terms in Eq. (1) can be given.

ANN training, i.e. the problem of determining the optimal number H of hidden nodes and the optimal numerical value for the synaptic weights in Eq. (1), is in practice solved as a non-linear optimization problem and has many empirical aspects. A database of input/output examples – for this application exclusively coming from either neutron irradiation experiments or nuclear power plants surveillance programs – is used as target for fitting. A good practice is to separate this database into a training and a reference set; the training set is only used to optimize the ANN, and the reference set is used to assess the accuracy of its predictions for new cases. The problem of how to define such sets is addressed in the next section.

2.1. Algorithms for the definition of training and reference sets

A good practice, before training the ANN, is to define training and reference sets from the available database. It is essential to make sure that both sets are equally representative of the domains in the input and output spaces, without overlap, in order to provide a rich training set that contains enough learning material. On the other hand, it is also important to keep enough pertinent and, most

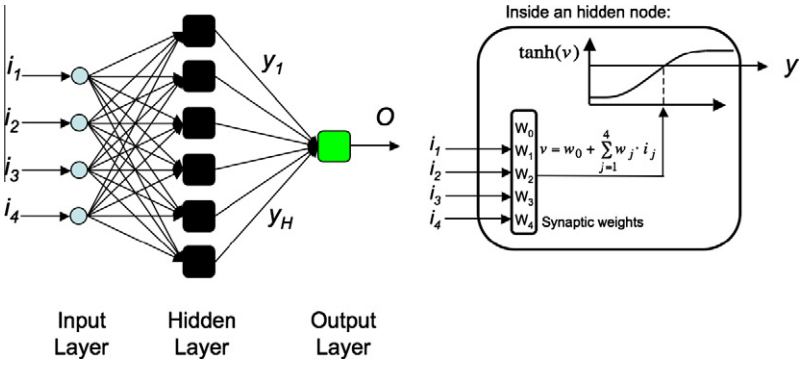

Fig. 1. Schematic representation of an artificial neural network (ANN) with four input variables (i_1–i_4), one hidden layer with six hidden nodes and one output (O). The input signals are propagated from left to right, in a layer-by-layer fashion, without feedback connection and without layer by-pass. The right part of the figure shows how a node processes its input signals into its own output.

importantly, independent examples for the validation of the trained ANN.

In some applications, the database can simply be split by "shuffling and cutting", like one would do with playing cards. This simple algorithm is, however, unsuitable for this application, for the following reasons:

(A) The distribution of points in the database input/output space is generally non-homogenous, as illustrated in Fig. 2. In particular, some regions are very sparsely populated, usually by just a couple of points. A completely random "shuffle-and-cut" algorithm could therefore lead to a poor sampling of these regions. Measures should be taken to avoid this.
(B) The "shuffle-and-cut" algorithm, even if biased to ensure proper representativity of both sets, inherently assumes that data points are sufficiently independent from each other, under the condition that at least one input variable assumes a different value. For this application, however, the input variables describing the chemical composition of a given steel are likely to operate in synergy and to have a dominant influence on the corresponding ANN output, $\Delta\sigma_Y$. It might thus be preferable to group the data points by steels and build training and reference sets operating on the data taken steel by steel, i.e. without separating data points referring to the same steel.

Based on these considerations, we have used two different algorithms to build training and reference sets from a given database:

- *Algorithm "by independent points"*: Here, we ignore the possibility of dominance of the input variables describing the chemical composition of a given steel and simply apply the "shuffle-and-cut" algorithm. However, we do not apply it in a completely blind (random) way. Instead, we impose the condition that all data points belonging to the sparsely populated regions of the input/output space (Fig. 2), should be equally distributed within the two sets.
- *Algorithm "by steel"*: The data points are grouped depending on the steel they belong to and these groups are then assigned to one or the other set, trying to share equally the distributions of chemical compositions. This is not straightforward because of the nature of the databases.

The former algorithm has the advantage of being simpler and of creating training and reference sets in which all input variables (and corresponding output values) are more equally distributed. The latter algorithm is of more convoluted application, but takes into account the fact that the examples used as reference for the ANN are steels and not independent concentrations of chemical elements.

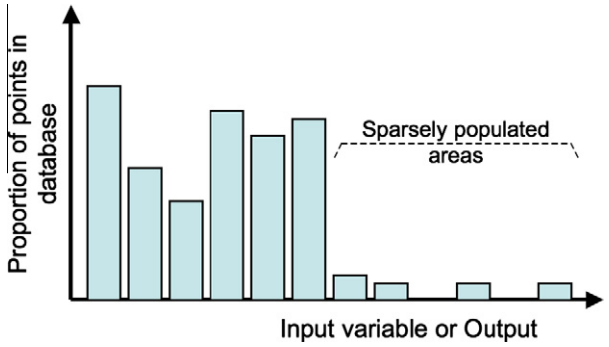

Fig. 2. Typical distribution of values for a given input variable in databases of irradiated steels. The variable in question can be for example the content of a certain chemical element or the irradiation temperature; the only output of interest for the RADAMO database is $\Delta\sigma_Y$.

2.2. Training algorithms

In this section, we describe the problem of training the ANN to reproduce as closely as possible the output data ($\Delta\sigma_Y$ in the present application), as a function of the input variables listed in the introduction.

We use the Levenberg–Marquardt (LM) training algorithm [13,14], *regularized* by *early stopping* [4]. The training set is used to update the synaptic weights in an iterative way, whereas the reference set is used to decide when training should be stopped, by assessing the actual ANN extrapolation skills on "never-seen" cases. In addition to the classical training scheme based on early stopping, we also considered the possibility of applying *node decay*, under a Bayesian training scheme [4], in order to compare the performance of the correspondingly obtained ANNs. The complete objective function f, to be minimized on the training set, is:

$$f = \beta \sum_{t=1}^{T} (d_t - O_t)^2 + \alpha \sum_{i=1}^{W} w_i^2 \tag{2}$$

Here, d_t is the desired ANN output for the training example number t out of T training examples and O_t is the corresponding network prediction. w_i is synapse i out of W. α and β are the *Bayesian hyperparameters*. The second term in Eq. (2) allows node decay to be introduced: it encourages the network to develop small value synapses connections, so as to yield the simplest possible regression. If this term is kept, the application of early stopping becomes theoretically unnecessary. In a classical LM training, with early stopping, we impose $\alpha = 0$ and $\beta = 1$, to turn off node decay. On the contrary, in a Bayesian training, α and β are not imposed, but are iteratively fitted. In this case, a variance on the ANN outputs can be theoretically calculated, to be later used as error bar:

$$\sigma^2 = \frac{1}{\beta} + g^T \cdot A^{-1} \cdot g \tag{3}$$

Here A is the matrix of the second order derivatives of the function f (see Eq. (2)) with respect to the synapses and g is the vector of the first order derivatives of the ANN output with respect to the synapses.

The reason for comparing the two training schemes is that the Bayesian scheme is expected to create a less complex and therefore more general ANN, presumably more suitable for extrapolation purposes. On the other hand, Bayesian trained networks generally commit larger errors on the reference set compared to networks trained with a classical algorithm, as a consequence of the right-hand side term of Eq. (2).

The input variables are linearly normalized between −0.150 and +0.150 as is common practice in ANN training. Flux and fluences, however, are linearly normalized on a logarithmic scale (i.e. we take their logarithms as input variables and normalize them between −0.150 and +0.150 afterwards), because they vary by several orders of magnitude in the range covered by the database used in this work. This normalization in log-space homogenizes well the distribution of flux and fluences in the database. The output is linearly normalized between −1 and +1, because the output of the ANN is bound to this range, due to the hyperbolic tangent activation function of the output node. The synaptic weights are randomly initialized between −1 and +1. Fig. 3 illustrates how the average error committed on the training and reference sets, after training, evolves with the number H of hidden nodes. The error on the training set always decreases with the addition of new nodes, because: (1) the function f is optimized

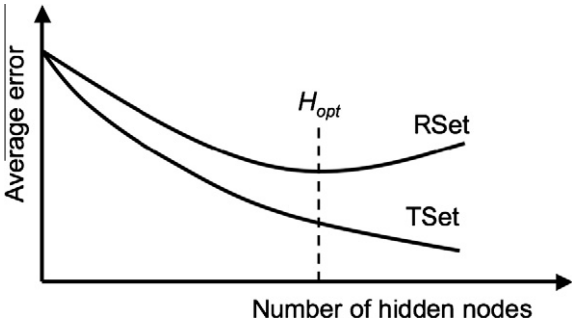

Fig. 3. Illustration of the usual dependence on the number of hidden nodes of the average error committed by the ANN, on the training set (TSet) and reference set (RSet). The dashed line shows the optimal number H_{opt} to retain, for the sake of generality.

on this set; (2) the ANN output derived from Eq. (2) is so general that the addition of new degrees of freedom helps improving the accuracy of the prediction. On the contrary, the error committed on the reference set stops decreasing after a certain optimal number H_{opt} of hidden nodes, because the ANN becomes more specialized for the list of examples used for training.

2.3. Use of network committees

A properly trained ANN is expected to provide good predictions for any new set of input variables, at least in the ranges covered by the training and reference sets. However, in practice no single ANN can be fully trusted because the mapping between input variables and output that is constructed during training is fitted "in a mean square sense". The network may have poorly learned the real effect of some input variables under particular sets of conditions that are insufficiently represented in the database. In other words, one should not expect an ANN to be able to predict physical phenomena that are not sufficiently represented in the database of examples used for training: the predictions made in partitions of the input space that are not properly covered by the training database are very likely to vary significantly from network to network, even if trained on the same database. For this reason, it is generally wiser to make predictions using a *committee* of networks, all trained on the basis of the same training and reference sets, rather than using individual networks. The final prediction of a committee of ANN, given a set of input variables, will be the average $\overline{O}$ of the predictions of the individual networks in the committee. In this way, a variance can be calculated:

$$\sigma^2 = \frac{1}{N}\sum_{i=1}^{N}\sigma^{(i)^2} + \frac{1}{N}\sum_{i=1}^{N}(O^{(i)} - \overline{O})^2 \quad (4)$$

Here, N is the number of networks in the committee, $\sigma^{(i)}$ is the standard deviation for the prediction of network i (calculated using Eq. (3)), and $O^{(i)}$ is the prediction of network i. Note that all $\sigma^{(i)}$ are 0 if the networks have been trained with the classical LM algorithm, without Bayesian node decay.

3. The RADAMO database

The RADAMO experimental program [15–18] was conducted at SCK•CEN to generate an experimental database covering a large spectrum of irradiation conditions. RADAMO was specifically oriented to measure irradiation effects on the tensile properties of RPV materials. Pressurized water reactor (PWR) and VVER[1] materials (plates, forgings and welds) with various chemical compositions were irradiated in the BR2 material test reactor under well controlled conditions at two temperatures, $T = 300\,^{\circ}C$ and $T = 265\,^{\circ}C$ in a large neutron fluence range from low ($\Phi < 10^{23}$ n/m^2) to high ($\Phi > 1.5 \times 10^{24}$ n/m^2, energy > 1 MeV) and various flux levels ($\varphi = 0.2$–8×10^{17} n/m^2/s, energy > 1 MeV).

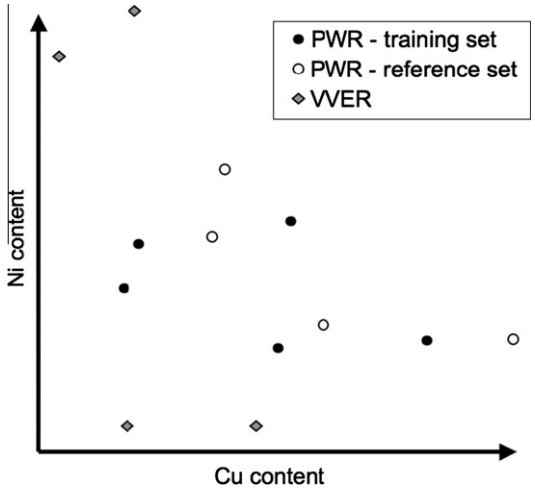

Fig. 4. Distribution of points in the Cu–Ni content space for the nine PWR steels and the four VVER steels contained in the RADAMO database. The PWR points are separated in training and reference sets, defined "by steels".

The database contains 346 entries related to PWR materials (nine different steels) and 63 entries related to VVER materials (four different steels). The input variables are: neutron flux, neutron fluence, irradiation temperature, chemical content with respect to several elements (Cu, Ni, P, Mn, Si, etc.) and also the product form (plate, forging or weld). The output is the variation $\Delta\sigma_Y$ of the yield stress measured by tensile tests at room temperature. Fig. 4 depicts the distribution of the 13 different steels in terms of Cu and Ni contents. We see that VVER materials are significantly different from PWR steels, especially in terms of Ni content. At this stage, in this work, we only use the PWR data to train the ANN. VVER data are kept in a separate database, to be used at a later stage to assess the performance of the trained ANNs, when employed to extrapolate to compositions significantly different from those represented in training and reference sets.

Before defining training and reference sets from the PWR data, we also separate a set corresponding to the highest neutron fluences: set PWR_HF, which contains 20 data points, all with neutron fluence higher than, or equal to 1.5×10^{24} n/m^2. As in the case of the VVER data, this set will be used at a later stage to test the performance of the trained ANNs, when employed to extrapolate to high fluence.

In conclusion, only the 326 remaining data points were used to build training and reference sets. These, as explained in Section 2.1, where defined in two different ways:

(A) *by independent points*, with 162 data points in the training set (PWR_TA) and 164 in the reference set (PWR_RA);
(B) *by steels*, with 174 data points in the training set (PWR_TB), corresponding to five different steels, and 152 in the reference set (PWR_RB), corresponding to four steels.

4. Results

4.1. Identification of the most influential input variables

In this section we describe the differences between ANN committees trained on different sets of input variables, in order to

[1] VVER is the abbreviation used to denote Russian-type light water reactors.

identify which are the dominant ones determining the output, $\Delta\sigma_Y$. Neutron fluence and irradiation temperature are unquestionably important and were therefore always included. The training experiments presented here concern the effect of the following variables: neutron flux, product form, and chemical composition.

For these experiments, we exclusively used the training and reference sets defined with the algorithm "by steel", in order to minimize the risk of specializing the ANN for all nine RPV steels of the database. The classical LM training algorithm is the most appropriate, because it allows more precision to be achieved within the given database, as will be clearly shown later.

In order to evaluate the performance of the ANN, we compared the predicted values with the values in the reference set. Better performance corresponds to better correlation between the two values, assessed using Pearson's product-moment correlation coefficient, R^2, as well as by calculating the average error, $\bar{e}$, defined as:

$$\bar{e} = \frac{1}{R}\sum_{r=1}^{R} |d_r - O_r| \tag{5}$$

Here d_r is the desired ANN output for the reference example r out of R, and O_r is the actual ANN output for the same input variables.

Extensive studies led us to the following conclusions:

- Product form and neutron flux are apparently variables that do not influence the output $\Delta\sigma_Y$, because, all other input variables being identical, their inclusion or exclusion does not change the ANN quality of prediction. Therefore, we conclude that no significant product form or flux effect is numerically discernable, at least in the range of irradiation conditions and chemical compositions covered by the RADAMO database.
- Amongst the chemical composition input variables, Cu content is unquestionably (and unsurprisingly) the dominant one. Fig. 5 shows that, even by considering only the Cu content as chemical composition variable, the correlation between predicted and reference data is very strong.
- The addition of the content of a second chemical element as input variable improved the correlation. However, the improvement achieved is almost the same independently of the choice of the second element (Ni, Mn and Si).

This last point is illustrated in Fig. 5, which shows ANN predictions obtained when neutron fluence, irradiation temperature and contents of two chemical elements are taken into account. We see first of all that the accuracy of the predictions is, in general, very high, and is, on average, of the same order as the experimental uncertainty. We also see that the difference in the quality of the prediction when considering Ni, Si, or Mn content as second chemical composition variable, is almost indiscernible. Taking any of them into account, in addition to Cu, hardly makes any difference in the final ANN accuracy.

We can suggest several explanations for this interesting result:

- Copper has a clear and distinguishable effect on the output $\Delta\sigma_Y$, which is very easily learned by the ANN. The other elements also have a distinguishable influence, separate from the copper effect, but there is a synergy between them, so that it is hard to isolate the individual effects of Ni, Mn and Si. So, the ANN understands the synergic effect without distinguishing the actual role of each element, possibly because of the limited number of steels compositions in the database.
- The ANN is not learning a general logic, but is in fact constructing a non-physical artifact that minimizes the function f in Eq. (2), by somehow "memorizing" the steels of the database, therefore making predictions that are only sound for them. This is however unlikely, because the algorithm "by steel" to define the training and reference sets is specifically aimed at avoiding this.
- The ANN manages to find a correlation between Ni, Mn and Si contents that is not immediately visible to us, but does exist. Such a correlation may exist and be found because the steels

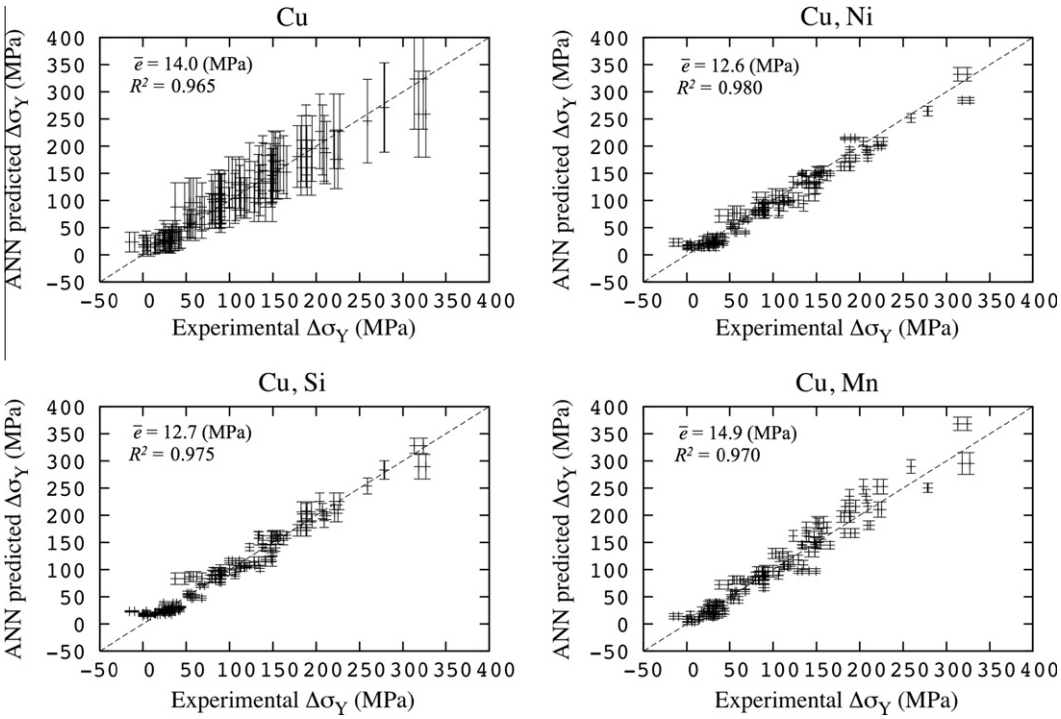

Fig. 5. Comparison of ANN predictions, measured on the reference set by committees of 30 networks (four hidden nodes), obtained when different series of input variables are taken into account. Training and reference sets were defined with the algorithm "by steel" and the ANN's were trained with the classical LM algorithm. In all cases, the input variables are: neutron fluence, irradiation temperature, and the content of the chemical elements indicated on the graphs. Error bars were calculated with Eq. (4), $\bar{e}$ was calculated with Eq. (5), and R^2 is Pearson's product-moment correlation coefficient.

of the database are not sampled in such a way that the chemical content of each element varies independently. On the contrary, the nine steels in the RADAMO database set can uniquely be labeled by just looking at the content of two chemical elements.

A way to differentiate between the importance of Ni, Si and Mn content as variables determining irradiation hardening is to investigate how the ANN is accurate for other steels, whose composition is far from the ranges covered by the training and reference sets. Fig. 6 shows the accuracy of the predictions on the VVER database, obtained by using the same committees of networks as in Fig. 5. We clearly see that: (1) the Cu content variable alone is not enough to ensure accurate extrapolations; (2) in this case Ni content appears the best one to retain, as the second most influential chemical element after Cu, and the one that provides the best extrapolation capabilities.

Finally, our attempt to improve the correlation by taking more than two chemical elements into account failed, because the ANN accuracy was never better than that shown in Figs. 5 and 6. This can either be explained by the fact that a hypothetical synergic effect of Si, Mn and Ni on hardening, plus the effect of Cu, removes the need to take all chemical elements into account, or, more simply, that the database does not contain enough training examples to correctly deal with such a regression problem, if more than four input variables are involved.

To summarize, we have shown in this section that the highest quality ANN predictions can be obtained by taking just four input variables into account: neutron fluence, irradiation temperature, Cu and Ni content. Neutron flux and product form have no significant influence on the RADAMO $\Delta\sigma_Y$ output.

4.2. Comparison between the proposed training schemes

In this section, we retain the four input variables that were selected based on the study reported in the previous section (neutron fluence, irradiation temperature, Cu and Ni contents) and compare the ANN accuracy of predictions after training with either the classical or the Bayesian LM algorithm, as well as defining training and reference sets either "by independent points" or "by steel" (see Section 2.1).

4.2.1. Comparison of the set definition algorithms

In Fig. 7, the performance of ANN committees trained with the classical LM algorithm is shown, defining the training and reference sets either "by independent points" (left side) or "by steel" (right side). In the same figure (lower part), the ANN committee is also tested on the high fluence set, PWR_HF. The biases $\bar{b}$, when shown, were calculated as:

$$\bar{b} = \frac{1}{N}\sum_{i=1}^{N}(O_i - d_i) \tag{6}$$

where d_i is the desired ANN output for the high neutron fluence example i out of N, and O_i is the actual ANN output for the same input variables. We see that the predictions on the reference sets (upper part of the figure) are slightly more accurate when the algorithm "by independent points" is used. The predictions on the high fluence set (lower part) are slightly more biased when the algorithm "by steel" is used. This can be explained by the fact that the ANN, in the first case, is more specialized for the particular steels represented in the database, and therefore manages to perform a more accurate extrapolation on the neutron fluence variable. As an additional illustration, the performance of the same ANN committees is illustrated, in Fig. 8, where predictions are obtained for the steels of the VVER database, i.e. with chemical compositions that are far from those used during training, either in the training or reference sets. We can see that the committee of networks yields slightly better extrapolations when training and reference sets were defined "by steel".

It therefore appears that the use of the algorithm "by independent points" to define the training and reference sets is preferable for the purpose of extrapolation on the neutron fluence variable,

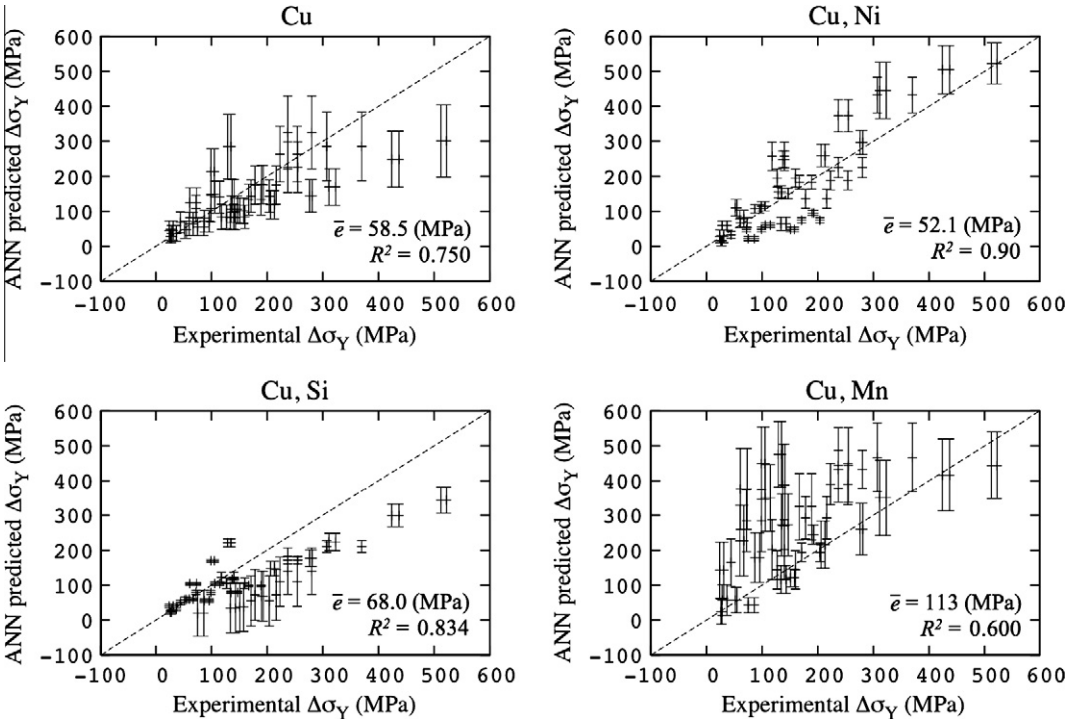

Fig. 6. ANN quality of prediction for the VVER database, using the same committees of networks as in Fig. 5, i.e. trained using the PWR data. (See Fig. 5 caption for the definition of $\bar{e}$ and R^2).

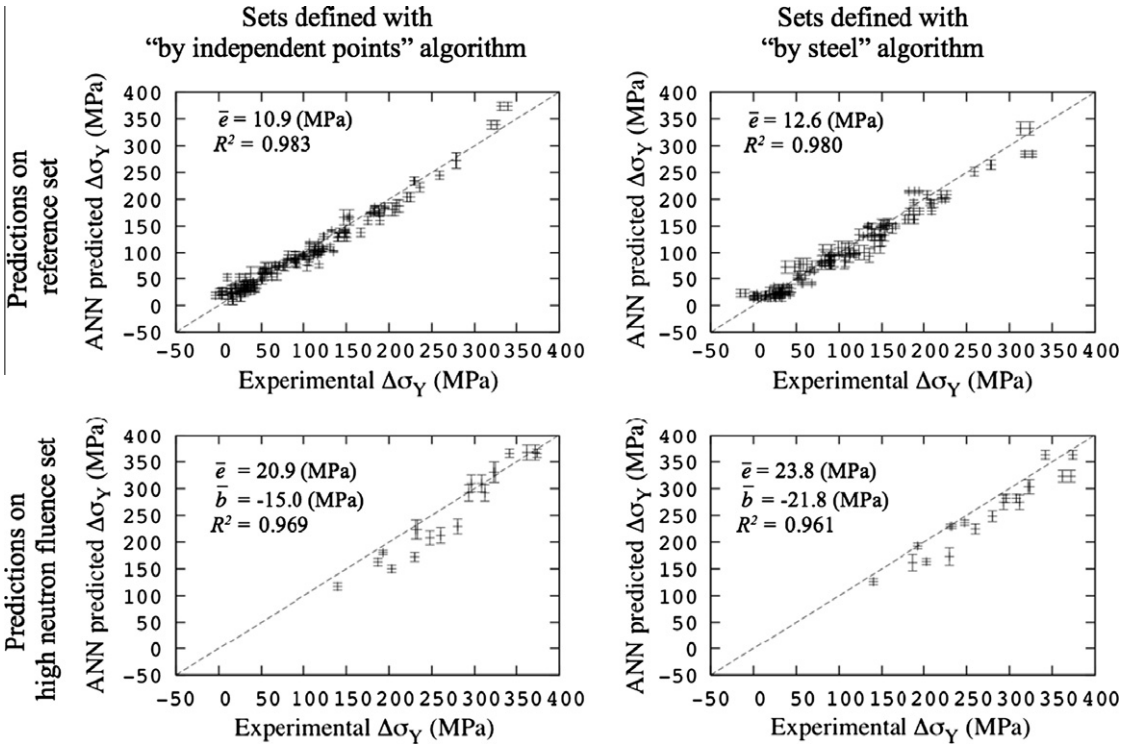

Fig. 7. Quality of prediction for committees of 30 networks (four hidden nodes), trained with the classical algorithm. Training and reference sets were either defined "by independent points" (left side) or "by steel" (right side). Error bars were calculated with Eq. (4), $\bar{e}$ with Eq. (5), $\bar{b}$ with Eq. (6), and R^2 is Pearson's product-moment correlation coefficient. Upper part: predictions for the reference set; lower part: predictions for the high neutron fluence set.

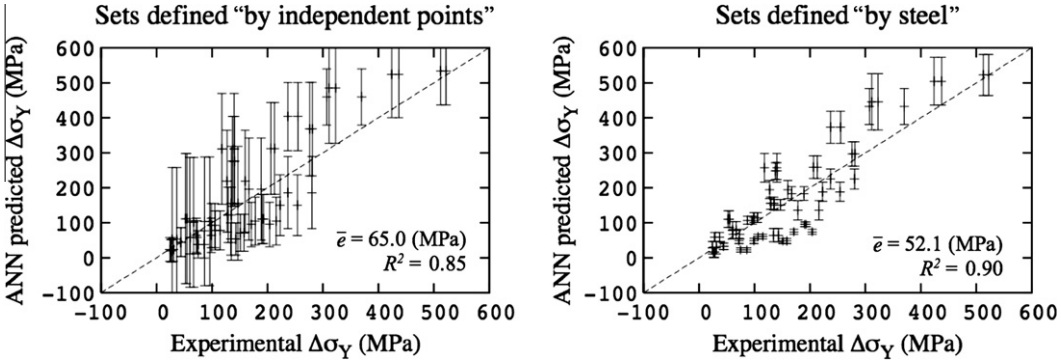

Fig. 8. ANN quality of prediction measured on the VVER database, using the same committees of networks as in Fig. 7, i.e. trained using the PWR data. (See Fig. 6 caption for the definition of $\bar{e}$ and R^2).

considering a steel that is already represented in the database. Otherwise, the use of the algorithm "by steel" is recommended, in order to predict the hardening of steels not included in the reference database.

4.2.2. Comparison between classical and Bayesian training algorithms

Fig. 9 is the equivalent of Fig. 7, for ANN committees trained with the Bayesian LM algorithm, defining training and reference sets either "by independent points" (right side) or "by steel" (left side). In comparison to Fig. 7, we see that the general accuracy of prediction is significantly lower than when classical LM training is used. In particular, the biases of the predictions on the high fluence set are much larger. This can be explained by the introduction of node decay in Eq. (2) (Section 2.2).

The predicted dependence of hardening on fluence for different temperatures, for two steels of the reference database, is shown in Fig. 10. We see that the Bayesian trained network is accurate for the higher temperature (300 °C), and is in fact closer to experimental data than the other network (trained in a classical way), because the average prediction line touches the bulk of experimental points and the error bands calculated with Eq. (4) encompass all of them. The quality of the predictions for the lower temperature (265 °C), however, is poor compared to the classically trained network. This can be explained by the fact that this temperature is poorly represented in the RADAMO database: only 48 data points, i.e. less than 15% of the database. As the minimization of the function f, in Eq. (2), is solved as a mean-square optimization problem, and as node decay prevents the network from developing a complex structure, the ANN became specialized for the higher temperature, that largely dominates the database.

One possible way to improve the generality of the Bayesian trained ANN, for the irradiation temperature variable, could

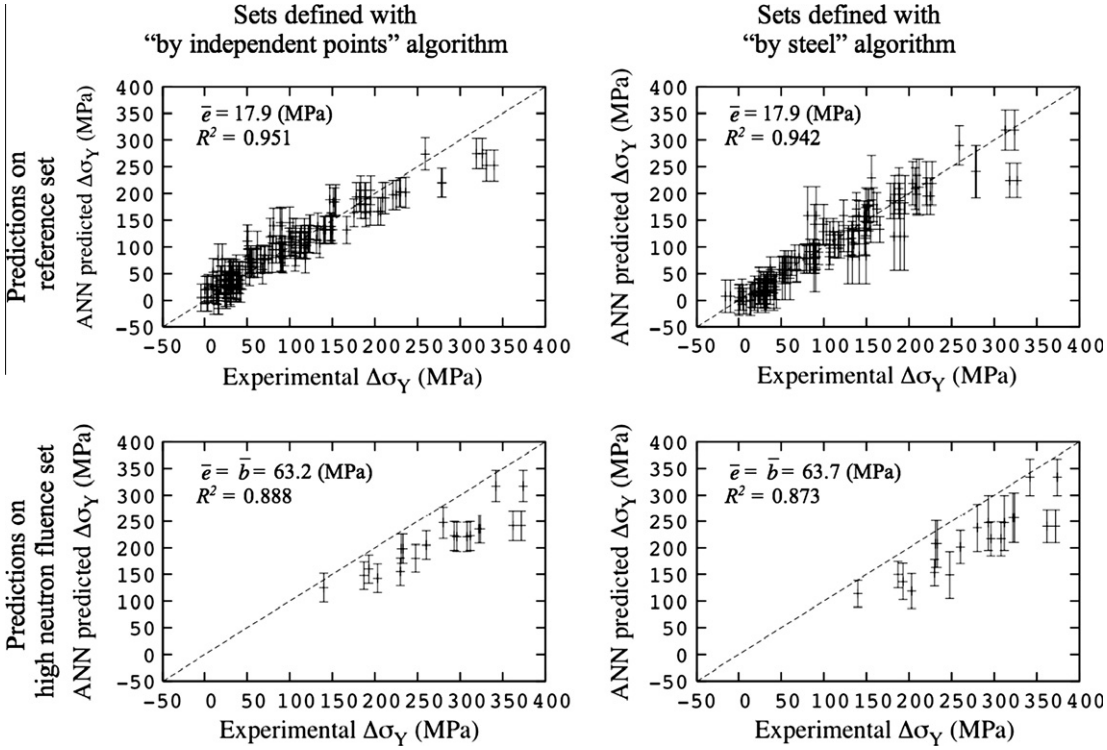

Fig. 9. Quality of prediction for committees of 30 networks (four hidden nodes), trained with the Bayesian algorithm. Training and reference sets were either defined "by independent points" (left side) or "by steel" (right side). Error bars were calculated with Eq. (4), $\bar{e}$ was calculated with Eq. (5), $\bar{b}$ was calculated with Eq. (6), and R^2 is Pearson's product-moment correlation coefficient. Top: predictions for the reference set; bottom: predictions for the high neutron fluence set.

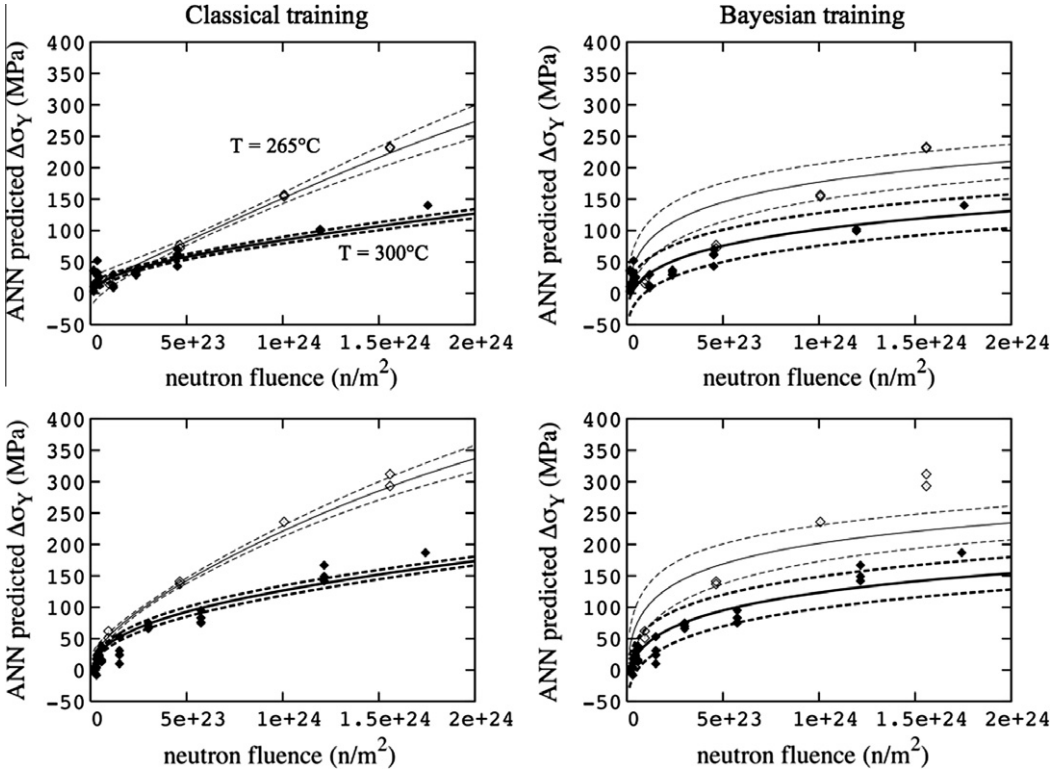

Fig. 10. Evolution with neutron fluence of the ANN-predicted $\Delta\sigma_Y$, versus experimental measurements, for two steels. Predictions are obtained using a committee of 30 networks (four hidden nodes). The training and reference sets were generated "on independent points". Dashed lines show the error bands calculated with Eq. (4). Left side: classical training; right side: Bayesian training.

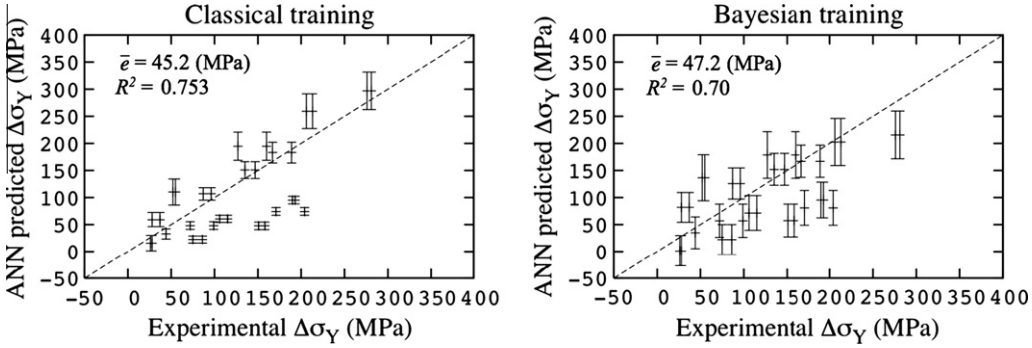

Fig. 11. Quality of prediction for committees of 30 networks (four hidden nodes), trained with the PWR data (training and reference sets were defined with algorithm B), for the VVER data (T = 300 °C). The training algorithm was either classical or Bayesian LM. Error bars were calculated with Eq. (4), $\bar{e}$ was calculated with Eq. (5), $\bar{b}$ was calculated with Eq. (6), and R^2 is Pearson's product-moment correlation coefficient.

therefore be to homogenize the database, by removing several points corresponding to 300 °C, so that each temperature (265 °C and 300 °C) represents about 50% of the database. We thus tried re-training with only 108 data points of the database: nine steels times two temperatures times six data points per steel and per temperature. Unfortunately, even after this homogenization the Bayesian trained ANN accuracy is not improved for both temperatures. This could be the consequence of the limited amount of steels in the database, which makes the application of Bayesian training inappropriate: node decay is apparently too strong a constraint that inhibits the ANN quality of prediction, despite its theoretically higher extrapolation skills. In other words, Bayesian training seems to express its full potential only if a sufficiently homogeneous *and large* database is available for training.

To conclude this study, Fig. 11 compares the accuracy of the predictions provided by the classically trained or by the Bayesian trained ANN committee on the VVER database. Only data points corresponding to 300 °C are shown, since we have already observed that the Bayesian network is inappropriate for the lower temperature. We can see that the accuracy, estimated on the basis of mean error and correlation, is very similar in both cases, although slightly lower for the Bayesian trained ANN committee. We may summarize that, if trained on a database larger than RADAMO, the accuracy achieved with Bayesian training would be higher than with classical training. However, in this specific case, very similar results are obtained in both cases when extrapolating to the different chemical compositions found in VVER steels.

5. Concluding remarks

In this paper, we have shown that artificial neural networks can be used to accurately predict neutron irradiation induced hardening of reactor pressure vessel steels, by taking into account the neutron fluence, irradiation temperature, Cu and Ni contents. The advantage of this numerical regression technique is that no hypothesis about how exactly these input variables influence hardening needs to be explicitly formulated, as the artificial intelligence approach itself takes care of finding non-evident relationships between the input variables and the output. This allowed us to point out the apparently negligible influence of the neutron flux, as well as of the product form, as input variables determining hardening, at least within the range of chemical composition and irradiation temperatures covered by the RADAMO database. At the same time, we have shown that the determination of the most influential chemical elements for hardening based on purely empirical considerations is not straightforward.

We have compared two training algorithms, as well as two methods for defining training and reference sets from the available database. We concluded that training the artificial neural network with an early stopping regularized algorithm, without the application of node decay, and defining training and reference sets "by independent points", can be recommended in order to train networks that can be accurately extrapolated from an existing database to high neutron fluences.

However, the accuracy in the extrapolation to different chemical compositions (VVER) is not fully satisfactory. In particular, the application of Bayesian node decay as a way to construct better networks provided no significant improvement, probably because of the limited amount of examples in the RADAMO database. In future work, we will further investigate the possibility of extrapolating to different chemical compositions, by extending the database and/or combining several databases of irradiated steels. The application to surveillance data will be our ultimate objective.

Acknowledgements

The authors are grateful to Dr. Enrico Lucon (National Institute of Standards and Technology, Boulder, USA) and Dr. Sehila Maria Gonzalez de Vicente (European Fusion Development Agreement, München, Germany) for their participation and interest in this work, fruitful discussions and comments on the paper.

References

[1] R. Gérard, Survey of National Regulatory Requirements, AMES Report No. 4, EUR 16305 EN, Brussels, June 1995.
[2] P. Petrequin, A Review of Formulas for Predicting Irradiation Embrittlement of Reactors Vessel Materials, AMES Report No. 6, November 1996.
[3] IAEA Nuclear Energy Series, Integrity of reactor pressure vessels in nuclear power plants: assessment of irradiation embrittlement effects in reactor pressure vessel steels, IAEA Nuclear Energy Series, No. NP-T-3.11, 2009.
[4] C.M. Bishop, Neural Networks for Pattern Recognition, Clarendon press, Oxford, 1995.
[5] F.F. Soulié, P. Gallinari, Industrial Applications of Neural Networks, World Scientific Publishing Co., 1998.
[6] J.R. Rabunal, J. Dorado, Artificial Neural Networks in Real-Life Applications, Idea Group Publishing, 2006.
[7] R. Kemp, G.A. Cottrell, H.K.D.H. Bhadeshia, G.R. Odette, T. Yamamoto, H. Kishimoto, J. Nucl. Mat. 348 (2006) 311–328.
[8] C.G. Windsor, G. Cottrell, R. Kemp, Modell. Simul. Mater. Sci. Eng. 16 (2008) 025005.
[9] S. Haykin, Neural Networks: A Comprehensive Foundation, MacMillian, New York, 1994.
[10] K. Hornik, M. Stinchcombe, H. White, Neural Networks 2 (1989) 359–366.
[11] A.-S. Bogaerts, R. Gérard, R. Chaouadi, Comparison of embrittlement trend curves to high fluence RPV surveillance results, in: International Symposium on Contribution of Materials Investigations to Improve the Safety and Performance of LWRs (Fontevraud 7), Avignon, France, 2010 (26–30 September).
[12] R. Chaouadi, R. Gérard, A.-S. Bogaerts, Assessment of the French and US embrittlement trend curves applied to RPV materials irradiated in the BR2 reactor, in: International Symposium on Contribution of Materials

Investigations to Improve the Safety and Performance of LWRs (Fontevraud 7), Avignon, France, 2010 (26–30 September).
[13] K. Levenberg, Q. Appl. Math. 2 (1944) 164.
[14] D. Marquardt, SIAM J. Appl. Math. 11 (1963) 431.
[15] R. Chaouadi, R. Gérard, J. Nucl. Mat. 345 (2005) 65–74.
[16] R. Chaouadi, An engineering radiation hardening model for RPV materials, Belgian Nuclear Research Centre (SCK•CEN) Report, R-4235, September, 2005.
[17] R. Chaouadi, RADAMO – An experimental databank for investigating irradiation strengthening of RPV materials, Belgian Nuclear Research Centre (SCK•CEN) Report, R-3858, April, 2004.
[18] R. Chaouadi, Irradiation-induced hardening of eight RPV steels and welds – RADAMO irradiation, Belgian Nuclear Research Centre (SCK•CEN) Report, R-3676, December, 2002.